

3.4 Increase in the Number of Failures Due to Cutting Soak Period in Half

We now estimate the increase in the number of outages caused by cutting the soak period. Reducing the soak duration results in some faults being passed on to the field that would otherwise have been previously found and corrected. We assume that the number of missed faults is proportional to the number of problems estimated to be found during the next 60 office-weeks following soak.

We used the Problems Found per Office Weeks for Switch G found in a supplier publication. We fit an exponential model to this data. The following table illustrates the results:

Table 10: Failures Due to Reduction of Soak Period

Description	Soak Time	Estimated Problems Found Seen in next 60 Office-weeks	Predicted Number Of Post-Soak Field Faults (Expressed As Ratio To Normal)
75% Reduction In Normal Soak	60 office weeks	132	$3.3 = 132/41$
50% Reduction	120 office weeks	83	$2.0 = 83/41$
25% Reduction	180 office weeks	52	$1.3 = 52/41$
Normal Soak Time	240 office weeks	41	$1 = 41/41$

The table says that if we cut the soak period by 75%, we would see about 3.3 times more troubles in the first 60 office-weeks than we would if the soak time was not decreased at all. Similarly, with a 50% reduction in the soak time, we would see about 2 times as many troubles in the first 60 office weeks than with the normal soak time.

4. Switch Failures Due to Increased Traffic

4.1 Discussion

LNP requires the originating switch to do several new things. These activities take time and switch resources. Call processing will be much more complicated with either the LRN or QoR approaches. The amount of additional call processing depends on the mix of traffic in an office and the percentage of numbers that are ported. However, it is clear that, as a result of the introduction of LNP, some switch processors will be upgraded and some switches may have to be reconfigured (i.e., split) to accommodate the increased load. We assume in this section that a switch overload will cause a severe performance degradation that customers may perceive as a switch outage.

For LRN, potentially every interswitch, intraLATA call will require a query to the LNP database. For QoR, only calls to a ported numbers will require a query to the LNP database. These queries take time and use resources of the local switch. For typical call mixes as specified in the IAO-25 model, we can calculate the effect on the LEC CCS network in terms of the number of octets that have to be processed. In the following table we indicate the multiplicative factor in the number of octets:

Table 11: Increase in CCS Network Traffic Due to LNP Introduction

	LRN	QoR
1% Ported	2.09	1.15
5% Ported	2.37	1.57
10% Ported	2.68	2.03
15% Ported	2.95	2.44
20% Ported	3.16	2.78

Preliminary information indicates that the introduction of LNP has a large potential impact on one of the widely used local switches, and that some processors used in these switches will need to be upgraded to avoid overloaded processor conditions. If the processors are not upgraded, some switches may be in real danger of overload when LRN is implemented. Information from the supplier indicates that LRN will require about a 30 - 47% increase in processing time in a typical switch. Any of these switches having a processor with a utilization that is greater than 50% before the introduction of LNP should be closely examined before LNP is introduced in that office to determine if a processor upgrade is necessary. For calls to portable NPA-NXXs (such that a query is performed at a switch, the call processing time increases up to 200% over the call processing time for a POTS call. With QoR, there is no increase in call processing time for calls to numbers that are not ported. However, calls to ported numbers result in an increase of 258 to 277% in call processing time at the switch that must perform the LNP query, and also result in about 50% of the processing time of a call that completes at the donor switch. These substantial call processing time increases in the switch can greatly increase the probability of a switch outage if they are not

taken into account in LNP planning. We theorize that the major reasons for this large effect on this type of switch are the centralized architecture that it uses and the specific AIN implementation.

Preliminary information for another widely used local switch is that it will experience less increase in call processing time than the switch described previously. The supplier for this switch has stated that the additional call processing time for each call will be about 10 to 30% greater than the POTS call processing time. We theorize that this is because of the specific switch's distributed architecture and AIN implementation.

The remaining widely used local switch also has centralized processing. This means that it also should see some of its capacity erode when LRN is implemented. The supplier has indicated that one of these switches making a query will experience about a 20% increase in call processing time over the time required to process a POTS call.

5. Simultaneous STP Failures

5.1 Discussion

With LNP, new software will be installed in each STP in the Houston Network. There are currently 2 STPs in Houston. The failure of one STP is transparent to customers. In order for customers to be affected both STPs must fail simultaneously. Note that the failure of both STPs will result in the loss of virtually all interoffice calls. It is a catastrophic outage. This section discusses the effects of putting in new software in the STPs on the chance of having both STPs fail during the 1st Quarter after cutover.

5.2 Background on STP Failures

Generally STPs are very reliable. In a Bellcore study of STPs, the probability that during a year an individual STP was out for over 30 minutes was 2.4%. The probability that an individual STP will experience an outage lasting longer than 30 minutes during a quarter is about 0.6%. All STPs come in mated pairs. For customers to experience an outage, both STPs have to fail simultaneously.

There are two majors ways that both STPs can fail:

- 1) we can have a common failure mechanism (e.g., common software) in each STP, or
- 2) a traffic overload brings both STPs down.

In this section, we will estimate the chances of a common failure mechanism in each STP. The following section discusses a traffic overload.

5.3 Both STPs Fail Due to a Common Failure Mechanism

Risk analyses done at Bellcore currently use a probability of 0.01 that an STP will fail given its mate has failed. The chance that an individual STP fails in a year longer than 30 minutes is about

2.4%. The probability that either STP will fail is 4.8% ($= 2 * 2.4\%$). This means that the frequency of simultaneous STP failures is:

$$\begin{aligned} E(\text{Both STPs fail during a year due to a common failure mechanism}) &= 0.048 * .01 \\ &= .00048 \end{aligned}$$

This means that during a typical year after LNP there is about a 0.048% chance of losing both STPs simultaneously due to common failure mechanism. There is a 0.012% chance of losing both STPs in a quarter assuming that LNP has been deployed for about a year. (This figure does not include the effects of putting in new software, rapid soak, or rapid deployment. We include these affects in the next subsections.)

5.4 Increase in STP Failures Due to New Software Version

There will be an increase in the number of STP failures due to having a new version of software released. We are collecting data on the software failure history for STPs. In lieu of that data, we expect that the number of software failures follows a similar pattern to other large software systems. We estimate that, just like switches, the number of failures in the 1st quarter after a generic is cutover will be 3 times the number for succeeding quarters. This will result in three times the expected number of times that both STPs will fail due to a common failure mechanism. This multiplier will be used for both LRN and QoR since both will require significant new software in STPs.

5.5 Increase in STP Failures Due to Short Soak Period

We expect that the soak period for the DSC STP will be shortened by about fifty percent. In an earlier section, we discussed the effect of a shortened soak interval on local switch failures. We believe that a similar pattern holds for STPs. This will result in doubling the number of failures during the 1st quarter after soak.

5.6 Additional STP Failure Considerations

If the STP handles load sharing for the LNP databases, there is a possibility that something may go wrong with the load sharing software and an outage may result. It is of great importance that the STP be exhaustively tested. We plan to keep investigating failures of STPs along with other types of failures.

It will be necessary to enter new global title translations in the STPs both for LNP number translation services and for the Message Relay Service associated with LIDB and CLASS queries. There is a finite probability of human error in entering these translations which can cause impaired service or a network failure. These translations will change as additional NPA-NXXs are opened to portability, and each change offers the opportunity for errors that can result in impaired service or network failure.

6. STP Failures Due to Increased Traffic

6.1 Discussion

We have started investigating the chance of STP failures increasing due to increased traffic as a result of implementing LNP in Houston. LNP with the LRN approach will require the STP to pass a query to and a response back from the LNP database for every local call to a portable NPA-NXX. Our preliminary finding is that the current STPs appear to be able to handle the additional load due to the introduction of LRN. QoR with a small percentage of ported numbers is even less likely to appreciably increase the chances of failure due to overload. For Message Relay Service, no one seems to know yet what the effects might be.

7. LNP Database Failures Due to New Software and Hardware

7.1 Discussion

For LNP, new databases will be put in place in the network, and SWB has indicated that it plans to use the Bellcore ISCP to perform this function. Because of this, the Bellcore ISCP will be used as the example LNP database in this section of the report when it is necessary to cite a specific implementation of an LNP database. Since the Bellcore ISCP release 5.1 will be used, both new hardware and new software will be needed.

This section discusses the effects of new software. LNP databases will be used in a load sharing mode so that the failure of one LNP database should be transparent to users. (Traffic and overload issues are discussed in the next section.) The primary way to have large numbers of customers lose service is to have all LNP databases down at the same time. There is a major difference between LRN and QoR in the effect of this type of failure.

With LRN virtually all interswitch, intraLATA calls will be affected if the LNP databases fail. With LRN, the failure of all LNP databases could be a catastrophic outage because all of Houston would be affected. We have assumed that a failure of all the LNP databases will cause all the intraLATA interoffice calls to experience a three second time out interval while awaiting an LNP database response, and that the simultaneous timing out of all these calls will cause overloads in the local switches. Following the time out interval, the switches will attempt to complete the calls using pre-LNP routing (to a default switch for the dialed NPA-NXX). This assumes that the caller holds on during the three second time out interval. The local switch manufacturers are aware of this problem and their solutions to the problem will be tested in the next few months. We are assuming that they are unsuccessful. In addition, it is unclear how customers will react to at least an extra 3 seconds in the call set-up time. There may be other undiscovered failure modes that have similar effects.

With QoR, only calls to ported numbers will be affected by LNP database failures, and there really is little chance of a catastrophic outage. On the other hand, even with QoR, the failure of all LNP databases will mean that calls from all stations to ported numbers cannot

complete, and will thus eventually result in an FCC reportable outage if the LNP databases are down long enough and the percentage of ported calls is high enough. In the first quarter of installation, there will probably be less than 1% numbers ported. This means that even if both LNP databases fail simultaneously, there will be about 10,000 calls affected during a 30 minute outage. With 10% numbers ported, about 100,000 calls will be affected and the outage will be FCC reportable. The remainder of this section is aimed at providing information on how we estimated the chances of all LNP databases failing.

7.2 Background on LNP Database Failures

As of February 10, 1997, there have been no dual ISCP failures. Since all services that are considered critical are put on a mated pair of ISCPs, for a critical failure to occur both ISCPs must simultaneously fail. This means there have been no FCC reportable outages or catastrophic failures due to ISCP failures.

On the other hand, there have been single failures for CCS, AIN and 800 databases. In 1996, outages for CCS, AIN and 800 databases occurred at a rate of about 0.22 per year, but most of these outages were due to scheduled events. If we ignore scheduled events then the - outage frequency per year goes down to 0.04 per year. We will use 0.04 as the frequency of failures per year for ISCPs (after the 1st month in service.)

There are two major ways that we can have all LNP databases fail simultaneously:

- 1) we can have a common failure mechanism (e.g., common software) in each ISCP.
- 2) a traffic overload brings all the ISCPs down

In the remainder of Section 7, we will estimate the chances of the first type of failure. Section 8 discusses overload considerations.

7.3 All LNP Databases Failing Due To a Common Failure Mechanism

In the last risk analysis done for SWB in 1995, the probability that an ISCP would fail given that its mate failed was estimated to be equal to .01. The frequency of single LNP database failures per year is estimated to equal 0.04. Since, for LRN, SWB will have 4 LNP databases in Houston, the expected frequency of single LNP database failures will be 0.16 per year. Also, we assume that (since all LNP databases are from the same supplier) that if the one pair of LNP databases fails that there is a .9 chance that the other pair of LNP databases will fail. The frequency that all LNP databases in Houston will fail simultaneously is:

$$\begin{aligned} E(\text{All 4 LNP databases fail due to common failure mechanism per year for LRN}) &= .16 * .01 * .9 \\ &= .00144 \end{aligned}$$

This means that during a typical year after LNP there is about a 0.144% chance of losing all LNP databases simultaneously due to common failure mechanism for LRN⁶. There is a 0.036% chance of losing all 4 LNP databases in a quarter assuming that LNP has been deployed for about a year. (This figure does not include the effects of putting in new software, rapid soak, or rapid deployment.)

For QoR, only 2 LNP databases will be needed in Houston. The expected frequency of both LNP databases failing due to common failure mechanism is:

$$\begin{aligned} E(\text{Both LNP databases fail due to common failure mechanisms per year for QoR}) &= .08 * .01 \\ &= .0008 \end{aligned}$$

For a quarter, there is about a 0.02% chance of losing both LNP databases simultaneously.

7.4 Increase in LNP Database Failures Due to New Software Version

There will be an increase in the number of LNP database failures due to having a new version of software released. We do not have any data on the software failure history for CCS, AIN or 800 databases. We expect that the number of software failures follows a similar pattern to other large software systems. We estimate that, just like switches, the number of failures in the 1st quarter after a generic is cutover will be 3 times the number for succeeding quarters. This means that there is a 0.11% chance ($= .00144 * 3/4$) of all LNP databases failing for LRN and 0.06% chance ($= .0008 * 3/4$) for QoR in the first quarter after cutover.

7.5 Increase in LNP Database Failures Due to Short Soak Period

For LRN, Version 5.1 of the ISCP will be needed. Currently, Version 5.1 of the ISCP will be available by the end of August, 1997. We have to assume that the soak period for Generic 5.1 could be quite short (about 1/4 of the normal soak period). In an earlier section, we discussed the effect of a shortened soak interval on local switch failures. We assume that a similar pattern holds for LNP databases in general and the ISCP in particular. This will result in adding about 200% to the number of failures during the 1st interval after soak⁷.

A different aspect is the tradeoff between the risks due to a short soak period and risks due to changing out a version of software that has live traffic on it. For QoR, Version 5.0 of the ISCP could be used. If Version 5.0 is used, the soak period can be extensive. If Version 5.0 is used, the shortness of the soak period will not be an issue. However, Version 5.1 will ultimately have to be introduced into Houston. Installing Version 5.1 at a later time will result in changing the ISCP with live traffic on it. We do not know the risks of this type of change. Because of the unknown (and possibly huge risks involved), Southwestern Bell is planning on introducing LNP with Version 5.1. If Version 5.1 of the ISCP is used for QoR, then there will also be an additional increase of 200% in the number of LNP database failures during the 1st quarter due to the short soak period.

⁶ Note that we are assuming that these LNP database failures are not independent. If these failures were independent the probability of simultaneous failure of all four LNP databases would be microscopic.

⁷ As the implementation schedule for the LNP databases becomes more clearly defined, the actual soak period may be longer than anticipated.

7.6 Increase in LNP Database Failures Due to Rapid Deployment After Soak

For LRN, there will be 4 LNP databases installed in 4Q97. This is considerably faster than the deployment schedule was for the ISCPs that SWB currently owns. For local switches we estimated that the rapid deployment of a new generic would increase failures by about 12%. LNP databases will in all likelihood see a similar short term increase.

8. All LNP Databases Failing Due to Overloads**8.1 Discussion**

Below are the results from a preliminary traffic analysis of the ISCP and some preliminary estimates of the chances that the ISCPs can handle the load in Houston. We assume that there will be four LNP ISCPs in Houston because of the expected query volume. Four ISCPs with Version 5.0 does not have a large enough query per second capacity to support the expected volume of queries during major overloads in Houston. For LRN, we believe that 4 ISCPs using V5.1 will be necessary.

8.2 Background

In general, there have been few ISCP failures. In particular, there have been few ISCP failures due to overload; although, one such incident has occurred in October, 1994. In that failure, the ISCP not only was unable to provide service, but it could not be accessed for diagnosis and repair because the overwhelming call processing took priority over all other functions such as craft initiated status and diagnostic queries.

The stated capacities for four ISCPs are given in following table:

Table 12: Supplier Provided ISCP Capacity for a Four ISCP Complex

	All 4 working	One ISCP Failure
ISCP V5.0	2000 qps	1500 qps
ISCP V5.1	5400 qps	4050 qps

For Houston, we expect about 600 qps during the Busy Hour possibly going up to 800 qps. Furthermore, there may be groups of seconds during some Busy Hours when the number of queries per second will be between 1.5 to 2 times this average rate. This may put the ISCPs into overload frequently when the LRN approach is used for the ISCP V5.0 particularly if one of the ISCPs is not performing correctly. If one ISCP V5.0 fails, there is a great possibility that the other three will not be able to handle the anticipated traffic load if the failure occurs during the Busy Hour. Any extreme overloads are virtually certain to result in catastrophic outages for ISCPs V5.0. In addition, there would be problems in transitioning from ISCP V5.0 to V5.1. Thus, we concur with SWB's decision to use V5.1 initially to provide LNP network capability.

For ISCP V5.1, we do not believe that normal traffic fluctuations will result in the ISCPs being overloaded. Except for extreme cases, we anticipate that maximal traffic seen by the ISCPs to be within 1.5 to 2 times the average query rate during the Busy Hour. That is, we expect that the traffic will be relatively smooth. This is because the ISCPs will be seeing traffic from many different sources. When you combine traffic from many independent sources, the result tends to look like Poisson traffic. We looked at traffic over 2 links for 2 weeks. We collected query counts for each 5 second interval. The number of queries per 5 second intervals was remarkably flat. The number of queries per 5 seconds in hours with a lot of traffic exceeded 1.5 the average rate occasionally but never exceeded 2 times the average rate. Note that this does not include a focused overload or an overload due to some common cause (usually a natural disaster - hurricane, ice storm). Whether these unusual severe overloads cause the ISCPs to fail is not clear. It certainly means that thorough testing of the overload controls in the ISCP V5.1 must be conducted prior to cutover.

Additional Issue: There is an additional issue that could influence the chance of a major failure in LNP. With LNP (and particularly with the LRN implementation), massive corruption of the translation table in a LNP database is possible. Because this table is changed on an ongoing basis, there is a real chance that all the LNP databases translation tables could be corrupted. If this happens, the calls to ported numbers and the message relay service (which supports CLASS and LIDB queries) may come to a halt. Even a minor corruption in the tables could cause invalid responses and result in the SSP using default routing. The potential failure modes associated with this scenario will require some investigation.

9. Additional Issues

9.1 Procedural Failures Due to Rapid Installation

Human Resources Problem For Managing The Implementation Of a Flash-Cut Of LNP. The implementation of LRN will require installation of new equipment, new software, and new switch and STP translations in all of the switches and STPs of the Houston MSA. This means that a large number of craft who do not have experience with the LNP software (since LNP is a new network capability) will be involved in the cutover that must be scheduled over a short (3 month) interval.

For example, with the LRN approach to LNP, SWB will need to install 4 ISCPs V5.1 in Houston. SWB has installed 4 other ISCPs but the time interval for the installation of those four ISCPs was stretched out over 2 years. In Houston, the 4 ISCPs will need to be installed in three months. This means that SWB's Operations resources will be stretched very thinly with LRN.

9.2 Failures From Message Relay Service - Message Looping

Message Relay Service is used to route CLASS or LIDB queries to the proper destination in an LNP environment where the proper destination for a message cannot be determined by examination of the NPA-NXX contained in the global title address field of a query message. The message relay service performs a full 10-digit global title address translation on a message if the

number has been ported, or indicates that a 6-digit translation is sufficient for non-porting numbers.

If Illinois Commerce Commission (ICC) LNP SCP requirements (Issue 0.95, September 4, 1996) are used for the message relay service, there is a potential problem with message looping. The Illinois Number Portability Administration Center (NPAC)/Service Management System (RSMS) sends out broadcast messages concerning newly ported numbers, but does not coordinate the entry of these new translations in STPs or LNP databases. Thus, if SWB enters a ported number translation in its LNP database message relay database and starts routing CLASS or LIDB queries to a network that has not entered the comparable translation in its message relay point, then the query will be routed back to the SWB network (which will retranslate the query and reroute it to the other carrier). Thus, a message looping situation results. There is no industry-agreed solution to combat the message looping problem for all internetwork scenarios using the Illinois LNP SCP requirements, but both DSC and the Bellcore ISCP developers claim to have developed proprietary solutions to detect that the problem exists. To the extent that a solution to the message looping problem is not found, not implemented, or does not work properly, there will be a negative effect on network reliability, with the possibility of overloading the message relay point (the LNP database) and/or the SWB STPs. Note that a mechanism for preventing the message relay looping problem using a translation type mapping is proposed in GR-2936-CORE, but will require internetwork cooperation to implement.

If a message looping situation develops, it is instructive to have some quantitative feel for how much CCS network and link capacity would be used. In order to make this calculation, it is necessary to make several assumptions. These are:

- CCS links operate at 56 kbps
- The CCS message is 120 octets in length (this appears to be a good assumption for CLASS)
- The STP transit time is 66 msec (the nominal maximum value for a 120 octet message during non-failure operation from GR-82-CORE)
- The LNP database message relay time is 350 msec.

The message looping path shown in Figure 1 is assumed. More complex paths are possible.




Figure 1: Message Looping Path

In this case the looping message traverses two LNP databases and four STPs. Thus the message takes approximately $2 \times 350 \text{ msec} + 4 \times 66 \text{ msec}$ or 964 msec to loop. A 120 octet message that loops in 0.964 seconds has an effective rate of 124.5 octets per second. Since the capacity of a 56 kbps CCS link is 7000 octets per second, a single looping CLASS query will consume 1.78% of the capacity of each of the 6 CCS links shown in Figure 1. To the extent that the transit times of the LNP databases and STPs are faster than the nominal values shown, the loss of link capacity will be even worse. With the Illinois approach, the looping situation can persist for several minutes even if everything is working properly.

9.3 Other Message Relay Reliability Concerns

In addition to message looping, the use of the message relay service introduces the requirement for several other capabilities into the SWB network. These capabilities must be carefully implemented to avoid creating the potential for degraded performance or failures.

9.3.1 New STP Translations and Different Translations in Different Pairs of STPs

Note, the comments in this subsection only apply if ICC SCP requirements are used. If requirements specified in GR-2936-CORE are agreed to by all the carriers in Houston and used there, this subsection does not apply.

The translations in STPs that route queries to the message relay point must be changed from the translations that are used today (to point to the MRP). Changing the existing translations will provide the opportunity to make errors and create the potential for failures for CLASS and LIDB queries.

If the Illinois approach to message relay service is used, then it is necessary to route LIDB queries to a different STP pair for final translation than the STP pair that sent the message to the LNP database. The need to define different translations in different STP pairs using the same translation type leads to a source for errors that can cause a failure of LIDB queries.

9.3.2 Lack of Subsystem Management Capabilities in ISCP

Note, the comments in this subsection only apply if ICC SCP requirements are used. If requirements specified in GR-2936-CORE are agreed to by all the carriers in Houston and used there, this subsection does not apply.

The ISCP does not track and manage subsystem status for the subsystems to which it relays queries. For CLASS this means that if the CLASS subsystem at a switch is not operational, the ISCP will continue to send queries to it that will fail. This leads to additional CCS network messages and additional load on a node (the switch) that is already experiencing a failure condition.

The lack of subsystem management capabilities in the ISCP also means that the ISCP cannot control loadsharing for LIDB (if duplicated LIDB databases are used), and that an STP must be used to do final global title translations for LIDB. This causes an additional STP global title translation to be done and provides another potential location for translation errors to be encountered.

9.4 New Operations Systems and Interfaces

Introduction of the message relay service will make it necessary to share translations data among carriers using a regional SMS. In addition, SWB expects to use a local SMS to load translations data into its LNP databases. The development, interfacing, and operation of these systems will require coordination among carriers, system developers, and the entity which is chosen to operate the regional system. Each of these operations provides the potential for errors and increases the likelihood of failures.

9.5 Default Routing

Use of the LRN approach has some built-in problems. First, there is no specification in the FCC order which deals with which network is to do the LNP query. In the Illinois SSP requirements, the N-1 network is assumed (but not required) to do the query. Application of this principle to Houston means that an IXC which is terminating traffic to the Houston LATA is assumed (but not required) to query the call and terminate the call to the correct network. However, there is no requirement in the FCC order for the IXC to query the call. As a result, some smaller IXCs may elect to just terminate the call to SWB, and force the SWB network to perform the query function. This is called default routing, and it will provide a source of LNP query traffic that may not be accounted for in the SWB LNP traffic planning.

There is also a high likelihood that other carriers will send traffic that has not been queried to SWB's network during times of failure in the other carrier's network. This will force SWB's CCS network and LNP databases to handle spikes of unanticipated traffic because of failures in other networks. There is a real possibility that these unexpected traffic spikes may lead to failures in the SWB CCS network or some of the network nodes that are attached to it.

9.6 Effects on 911 Service

In addition to effects on switching and call setup that would affect 911 calls along with other calls, specific impacts on 911 service that have been identified to date include:

- Need for 911 tandem to serve more NPA-NXX codes. Depending on local conditions, the introduction of LNP may cause a substantial increase in the number of area codes that are needed in an area. This typically happens in areas with many rate centers or districts. If additional area codes are needed, they may impact the 911 tandem because of signaling methods that are used between the 911 tandem and the subtending end offices.
- Need for administrative flow from new carriers to input data to the E911 database. Information concerning ported and new native numbers in the E911 database. To the extent that updates concerning new networks numbers are not properly entered or updated in the E911 database, improper transfers to local public service answering points (PSAPs) or dispatches may be made.

9.7 Additional Traffic Issues

As indicated in Section 4, the introduction of LNP will cause a substantial step-function increase in the CCS network traffic - particularly if the N-1 network query approach is used. In this case, even if no numbers are actually ported, the CCS network traffic that results from call setup, typically the majority of CCS network traffic, will almost double because of the introduction of LNP (unless a query reduction mechanism such as the Query on Release technique is used). In the past CCS network traffic has increased gradually as additional nodes were added to the network. Subjecting the CCS network in Houston to this rapid increase in traffic load with inadequate testing provides one of the ingredients for a recipe for significant CCS network problems.

The only parallels that Bellcore has seen for this magnitude increase has been in countries where external events (such as natural disasters or bombings) led to very large, short term increases in CCS network traffic. The result in these countries has been both switch and CCS network failures lasting several hours as the result of the large, sharp increase in traffic. In those areas, the traffic increase was short lived. In the case of LNP, the increase is permanent and must be handled on an ongoing basis.

INFRASTRUCTURE

THE BUSINESS & TECHNOLOGY OF MULTIMEDIA NETWORKS

Hayes Microcomputer Products Inc. has licensed Asymmetric Digital Subscriber Line or ADSL technology from **Aware Inc.** and intends to develop ADSL modems for the consumer market. The licensing agreement covers Aware's Discrete Multitone, or DMT, technology and will enable Hayes to create modems that work with a telephone company's DMT-based service.

Bell Labs is moving in new directions, and creating an Internet-based map for the rest of us. **Lucent Technologies Inc.** announced its former research lab has created a new Web site, Maps on Us (www.MapsOnUs.com), that allows users to plan travel routes and search for U.S. locations on the Internet. The site combines yellow pages directories with directions and mapping features.

Ameritech Corp. won a five-year deal to manage **UAL Corp.**'s national voice network. The contract means the Ameritech GlobalDesk will become a telecommunications department to about 500 locations nationwide of UAL, the parent company of United Airlines. The multimillion-dollar deal is a major win for Ameritech.

Servis Communications, formerly known as Diamond Net, will use **Cascade Communications Corp.** equipment to build a national Asynchronous Transfer Mode, or ATM, network to provide broadband Internet and Intranet service. Servis is buying both frame relay and ATM switches from Cascade as it builds a network at 45 megabits per second, with a planned upgrade to 1.55 Mbps within the year.

Does Competition Mean Tie-Ups?

By Carol Wilson

A small but growing coterie of hardware and software vendors and industry consultants said the technical challenges of seamlessly connecting multiple networks may prove more difficult and costly than either the telephone companies or regulators believe.

As a result, widespread local service competition may not be at hand, a fact that could also delay Bell company entry into the long-distance market.

Of particular concern is the public network's signaling system — a separate data network that starts and stops telephone calls and enables features such as call forwarding, caller ID and automatic callback. Known as Signaling System 7, or SS7, this data network will also be required to deliver Local Number Portability, or LNP, a critical part of the local competition package. The incumbent telephone companies must provide LNP so consumers can change telephone companies without changing telephone numbers.

But the demands that LNP would put on the SS7 network will likely exceed its capacity, and the result could be disaster for the entire phone network. With number portability, every call made to a number that has been changed from one carrier's customer list to another's will require the signaling system to trigger a database "dip" to determine how to route the call.


Implementing LNP across the country will trigger an enormous number of such "dips," each consuming network resources. There is potential for the entire network to screech to a halt.

"You just have to look at the numbers," said Lee Smith, senior marketing director for Tekelec Corp.'s Switching Systems Division, which makes SS7 equipment. "The equipment today is engineered to handle 500 to 1,000 transactions a second. In a metropolitan area, LNP is likely to generate 10,000 transactions a second. That would require 20

times as much equipment as is in the network today."

Keith Willets, longtime operations guru for British Telecommunications PLC and founder of the Network Management Forum, now is senior vice president at TCSI Corp., a provider of

to happen," Willets said. "But this is your competitor's traffic and it isn't revenue-generating. The question then becomes, who is going to pay for all of that engineering work? The regulators in the U.S. and the U.K. have laid the lion's share of the burden on the



network management software and consulting services. In that capacity, he sees problems for SS7 networks.

"It's not just that you have a lot of traffic suddenly congesting your network, although that is absolutely going

to happen. This is going to be a serious problem."

While both state and federal regulators have looked at LNP and mandated solutions for providing the capability, there haven't been any major alarms set off on the issue of signaling network capacity.


"I haven't heard of this at all," said Terrence Barnich, former chairman of the Illinois Commerce Commission and currently a consultant in the local competition arena for New Paradigm Inc. "But it doesn't surprise me a bit. I'm seeing a lot of anecdotal evidence that the [local telephone companies] aren't prepared for competition in many arenas. There are just a lot of things they haven't considered yet."

Not everyone is alarmed, however. Ron Hoffman, special regulatory features manager for Lucent Technologies Inc., said SS7 technology will evolve over the next several years, beginning with an increase in link speed from 56 kilobits per second to 1.5 megabits per second.

While Federal Communications Commission rules on LNP have made life harder for incumbent carriers, the competition is everywhere at once.

"It's going to be a long, painful process," Hoffman said.

Continued on page 12



Bells May Not Be So Prepared To Compete

▼Continued from page I-1

LNP also could disrupt the telephone companies' internal systems, Smith said.

The accepted approach revolves around having a Numbering Plan

Administration Center within each service area, which would take service activations from incoming competitors and match them with service deactivations from the local incumbent to set up number portability. But

there are no defined interfaces between the center and the legacy Operations Support Systems that telephone companies use to keep track of network facilities.

"So when a service deactivation

order comes through, there is no direct link to be able to automatically update the company's legacy systems. That information would have to be entered separately, which is duplicated effort," Smith said.

Further complicating the situation is the Bell companies' spin-off of Bellcore, the research arm that developed most of the legacy systems. It likely would also have developed an industrywide LNP approach, then specified to vendors what would be required for implementation.

The LNP issue is just one part of the greater problem the industry will have to face, said Probe Research Inc. consultant Mike Arellano.

"At this point, there is no interconnection order, since it's being challenged in court, and there are a lot of other issues that have to be resolved, including who's going to pay for local number portability," he said. The bottom line, he said, is that local competition is nowhere near becoming reality. ▲



The Eagle, Proud and Strong

The Eagle has the strength and flexibility to conquer the rigorous demands of its environment.

Tekelec's EAGLE STP features open distributed architecture as well as high capacity and throughput. It is tailored to the SS7 switching needs of independent telcos, cellular carriers, IXCs, CAPs and RBOCs.

Installed in over 50% of domestic cellular companies that have their own STP's, the EAGLE is scalable from two to 700 links and has the capability to grow with your needs.

In a market with few certainties, the EAGLE is clearly the right choice at the right time.

The Tekelec EAGLE STP, Proud and Strong

For more information call Tekelec's
Network Switching Division at
Morrisville, North Carolina
919.460.5500 or visit us
on the internet at
<http://www.tekelec.com>



TEKELEC

More Than Technology...Expertise

See us at CTIA Wireless, booth
171, March 3-5 in San Francisco

Editorial Index

Companies	Pages
AT&T Corp.	I-1, I-12
AT&T Corp.	I-5, I-6
Averco Inc.	I-1
Bay Networks Inc.	I-14
Bell Atlantic Corp.	I-12
Bell Labs	I-1
BellSouth Corp.	I-5, I-12
British Telecommunications PLC	I-1
Cablevision Systems Corp.	I-7, I-10
Cablevision Systems Corp.	I-1
Carinet Services Inc.	I-7
Cisco Systems Inc.	I-14
Com21 Inc.	I-14
Comcast Cable Corp.	I-5, I-7
Continental Cablevision Corp.	I-7
Corning Inc.	I-7
Cox Communications Inc.	I-5, I-7
Cybernetic Output Inc.	I-10
First Virtual Corp.	I-10
GTE Corp.	I-12
Hewlett-Packard Products	I-7
Hewlett-Packard Co.	I-10, I-14
IBM Corp.	I-7, I-10
Krellco Systems Inc.	I-10
LAMCity Corp.	I-14
Lucant Technologies Inc.	I-1, I-14
NPS Communications Co. Inc.	I-7, I-12
NetScout Inc.	I-14
New Paradigm Inc.	I-1
Networks Networks Corp.	I-7
NextLevel Communications Inc.	I-10
NextWave Telecom Inc.	I-5
Nymex Corp.	I-12
Pacific Bell	I-6, I-12
Packard Bell Co.	I-6
Pentel Systems Corp.	I-6
PrimeCo Personal Communications LP	I-5
Qualcomm Inc.	I-6
RAMWorks Inc.	I-10
Rodgers Communications Ltd.	I-7, I-10
Savvis Communications	I-1
SBC Communications Inc.	I-12
Savvis Communications Co.	I-5
TCSI Corp.	I-7
Telecel Corp.	I-1
Tele-Communications Inc.	I-5, I-7
Teleport Communications Group Inc.	I-7
Time Warner Inc.	I-0
U.S. West Inc.	I-7, I-12
UAL Corp.	I-1
Unimod Planet Inc.	I-6
West End Systems Inc.	I-7

Advertisers Index

3Com	I-9
Adtran	I-1
Applied Innovations	I-2
General Datacomm	I-7
Holmco	I-4
Lucant Technologies	I-2, I-14
Paragon Technologies	I-18
Savvis Communications	I-5
Telecel	I-16